Catalogue no. 12-001-X
ISSN 1492-0921

Survey Methodology

Two local diagnostics to evaluate the efficiency of the empirical best predictor under the Fay-Herriot model

by Éric Lesage, Jean-François Beaumont and Cynthia Bocci

Release date: January 6, 2022

Statistics Canada Statistique Canada

Canada

How to obtain more information

For information about this product or the wide range of services and data available from Statistics Canada, visit our website, www.statcan.gc.ca.

You can also contact us by

Email at infostats@statcan.gc.ca

Telephone, from Monday to Friday, 8:30 a.m. to 4:30 p.m., at the following numbers:

• Statistical Information Service	1-800-263-1136
• National telecommunications device for the hearing impaired	1-800-363-7629
• Fax line	1-514-283-9350

Depository Services Program

• Inquiries line	1-800-635-7943
• Fax line	1-800-565-7757

Standards of service to the public

Statistics Canada is committed to serving its clients in a prompt, reliable and courteous manner. To this end, Statistics Canada has developed standards of service that its employees observe. To obtain a copy of these service standards, please contact Statistics Canada toll-free at 1-800-263-1136. The service standards are also published on www.statcan.gc.ca under "Contact us" > "Standards of service to the public."

Note of appreciation

Canada owes the success of its statistical system to a long-standing partnership between Statistics Canada, the citizens of Canada, its businesses, governments and other institutions. Accurate and timely statistical information could not be produced without their continued co-operation and goodwill.

Published by authority of the Minister responsible for Statistics Canada

An HTML version is also available.

Cette publication est aussi disponible en français.

Survey Methodology, December 2021
Vol. 47, No. 2, pp. 279-297
Statistics Canada, Catalogue No. 12-001-X

Two local diagnostics to evaluate the efficiency of the empirical best predictor under the Fay-Herriot model

Éric Lesage, Jean-François Beaumont and Cynthia Bocci[1]

Abstract

The Fay-Herriot model is often used to produce small area estimates. These estimates are generally more efficient than standard direct estimates. In order to evaluate the efficiency gains obtained by small area estimation methods, model mean square error estimates are usually produced. However, these estimates do not reflect all the peculiarities of a given domain (or area) because model mean square errors integrate out the local effects. An alternative is to estimate the design mean square error of small area estimators, which is often more attractive from a user point of view. However, it is known that design mean square error estimates can be very unstable, especially for domains with few sampled units. In this paper, we propose two local diagnostics that aim to choose between the empirical best predictor and the direct estimator for a particular domain. We first find an interval for the local effect such that the best predictor is more efficient under the design than the direct estimator. Then, we consider two different approaches to assess whether it is plausible that the local effect falls in this interval. We evaluate our diagnostics using a simulation study. Our preliminary results indicate that our diagnostics are effective for choosing between the empirical best predictor and the direct estimator.

Key Words: Empirical best predictor; Design mean square error; Model mean square error; Local diagnostic; Local effect; Fay-Herriot model.

1. Introduction

Governments need socioeconomic information at increasingly fine levels of detail. National statistical offices are therefore required to produce statistics for sub-populations that were not identified or could not be taken into account when the survey's precision objectives were determined. As a result, the number of sampled units for these sub-populations may be too small to ensure good precision of standard design-based direct estimators such as the Horvitz-Thompson estimator or calibration estimators. This type of sub-population, where the sample size is insufficient, is called a small domain (or small area). To remedy the lack of precision of direct estimators for small domains, indirect estimators, or small area estimators, can be used. These small area estimators usually rely on a model such as the Fay-Herriot model (Fay and Herriot, 1979). The Empirical Best (EB) predictor, also called the Empirical Bayes predictor or EB estimator, is a small area estimator frequently used in practice.

Small area estimation methods use statistical models to leverage information from the survey and from auxiliary data sources. The Fay-Herriot model is a linear model that breaks down the parameter of interest of a domain into two terms: the first term is the effect explained by the model and the second term is the model error that can be interpreted as an unexplained and unknown local effect.

Classical statistical tools, such as graphs of model residuals, can be used to assess the validity of the Fay-Herriot model. However, these tools give little indication of the efficiency of an indirect estimate for a particular domain. The model Mean Square Error (MSE) of an indirect estimator can be viewed as a

1. Éric Lesage, Insee, Regional Directorate, 35, place du Colombier - CS 94439 - 35044 Rennes Cedex. E-mail: eric.lesage@insee.fr; Jean-François Beaumont, Statistics Canada, R.H. Coats Building, 100 Tunney's Pasture Driveway, Ottawa, Ontario, Canada, K1A 0T6. E-mail: jean-francois.beaumont@statcan.gc.ca; Cynthia Bocci, Statistics Canada, R.H. Coats Building, 100 Tunney's Pasture Driveway, Ottawa, Ontario, Canada, K1A 0T6. E-mail: cynthia.bocci@statcan.gc.ca.

local quality indicator since it varies across domains. The model MSE accounts for the effect explained by the model, but integrates out the unexplained local effect.

The design MSE is an alternative to the model MSE that does not integrate out the unexplained local effect. However, unbiased design MSE estimates of small area estimators tend to be very unstable, particularly for domains with few sampled units (Rivest and Belmonte, 2000; Rao, Rubin-Bleuer and Estevao, 2018; and Pfeffermann and Ben-Hur, 2019). To circumvent this problem, literature suggests taking an average over several domains of the design MSE (Rao and Molina, 2015; and Pfeffermann and Ben-Hur, 2019) as a quality measure. However, many public statistics users are only concerned with their specific domain and do not buy into an overall quality criterion to assess the efficiency of estimates for their domain of interest. This is especially the case when they are convinced that their domain is very specific and that this specificity is not found in the explanatory term of the model, but rather in the error term, i.e., the unexplained local effect.

To address the problem of the instability of unbiased estimators of the design MSE, Rao, Rubin-Bleuer and Estevao (2018) proposed a composite estimator that they evaluated in a simulation study. Their composite estimator consists of taking a weighted average of a model MSE estimator and a design MSE estimator. They achieve greater stability at the cost of an increase in bias. Pfeffermann and Ben-Hur (2019) also proposed a method for estimating the design MSE of a small area estimator. The method is rather complex and relies mainly on the choice of an appropriate model. It is therefore not entirely design-based. Apart from these attempts to estimate the design MSE, to the best of our knowledge, there is no local diagnostic in the literature that can be used to determine whether small area estimation is preferable to direct estimation for a specific domain.

In this paper, a different approach is proposed to compare the efficiency under the design of the EB and direct estimators. We proceed in two steps. First, we determine the unexplained local effect interval that ensures the design MSE of the Best (B) predictor, also called Bayes predictor or B estimator, is smaller than the design MSE of the direct estimator. The second step is to assess whether it is plausible that the unexplained local effect lies within this interval. To this end, two diagnostics are proposed: one based on the conditional distribution of the unexplained local effect given the direct estimate, and a second based on a hypothesis test on the unexplained local effect carried out with respect to the sampling design. We found that, depending on the magnitude of the standardized model residual and a factor associated with the precision of the direct estimate, it is possible to detect whether the B or EB estimators are likely to have a smaller design MSE than that of the direct estimator.

Section 2 presents the Fay-Herriot model and describes how the best predictor (B estimator) of the population parameter of interest is constructed. In Section 3, the model and design MSEs of the direct estimator and best predictor are derived. Section 4 describes the two proposed diagnostics. Section 5 explains how to estimate the model parameters and obtain the empirical best predictor (EB estimator) and the estimators of the diagnostics. Section 6 presents the results of a simulation study using real auxiliary data. A brief conclusion is provided in Section 7.

2. The Fay-Herriot model and the best predictor

We consider a finite population U of size N and a sample s of size n drawn from U according to a sampling design $p(s)$. The population U is partitioned into m domains that do not overlap. The domains are identified by the subscript i taking values from 1 to m. The population of domain i, with a size of N_i, is denoted as U_i. The sample of domain i is denoted as s_i and its size is n_i. We are interested in estimating m finite population parameters, $\theta_i, i = 1, \ldots, m$, associated with the m domains. The parameter θ_i is usually a total, an average or a ratio for domain i. Auxiliary information is available in the form of vectors, $\mathbf{z}_i$, available for all domains $i = 1, \ldots, m$. The set containing the m auxiliary vectors is denoted by $\mathcal{Z} = \{\mathbf{z}_i\}_{i=1,\ldots,m}$. Furthermore, we denote by Ω, the set of all variables used to make inferences excluding the inclusion indicators in the sample s; Ω includes $\mathcal{Z}$ and $\theta_i, i = 1, \ldots, m$, among others. The design expectation of a random variable, say A, will thus be denoted by $\mathbb{E}(A \mid \Omega)$.

We consider a linking model that breaks down the parameters of interest θ_i as follows:

$$\theta_i = \boldsymbol{\beta}^\top \mathbf{z}_i + b_i v_i, \qquad i = 1, \ldots, m, \tag{2.1}$$

where $\boldsymbol{\beta}$ is a vector of model parameters of the same dimension as $\mathbf{z}_i$, b_i are fixed factors that can be used to account for heterosedasticity in the model and v_i are error terms that follow the normal distribution: $v_i \mid \mathcal{Z} \sim \mathcal{N}(0, \sigma_v^2)$, where σ_v^2 is a model parameter. In practice, $b_i = 1$ is a common choice but it may be more natural to choose $b_i = N_i$ when θ_i is a total. The term $\boldsymbol{\beta}^\top \mathbf{z}_i$ is the known effect or effect explained by the model of the finite population parameter θ_i, while $b_i v_i$ is the unknown or unexplained effect that is called the unexplained local effect of θ_i or simply the local effect of θ_i.

The direct estimator of θ_i is denoted by $\hat{\theta}_i$. It is usually obtained by assigning a survey weight to each unit of the sample s_i. The survey weight of a unit can simply be the inverse of its probability of selection in the sample s or a calibration weight. The sampling error is defined as:

$$e_i = \hat{\theta}_i - \theta_i. \tag{2.2}$$

In what follows, the direct estimator will be assumed to be design-unbiased, i.e. $\mathbb{E}(\hat{\theta}_i \mid \Omega) = \theta_i$ or $\mathbb{E}(e_i \mid \Omega) = 0$. This assumption is not always satisfied in practice, for example when using calibration weights, but we will make the usual assumption that the bias remains negligible. We will also assume that the direct estimator $\hat{\theta}_i$, and thus the error e_i, follows a normal distribution. As discussed in Rao and Molina (2015, page 77), the normality assumption of the errors e_i is possibly weaker than the normality assumption of the errors v_i because of the effect of the central limit theorem on $\hat{\theta}_i$. Of course, this effect is less pronounced for smaller domains. Under these assumptions, we have: $e_i \mid \Omega \sim \mathcal{N}(0, \psi_i)$, where $\psi_i = \mathbb{V}(\hat{\theta}_i \mid \Omega)$ is the design variance of $\hat{\theta}_i$. The sample size n_i can be very small, which can lead to poor precision of the direct estimator $\hat{\theta}_i$. This problem has been at the origin of small area estimation research.

By combining the model (2.1) and the expression (2.2), we obtain the combined model, also called the Fay-Herriot model:

$$\hat{\theta}_i = \boldsymbol{\beta}^\top \mathbf{z}_i + b_i v_i + e_i. \tag{2.3}$$

Noting that v_i is fixed under the sampling design, it can easily be shown that $\mathbb{V}(b_i v_i + e_i \mid \mathcal{Z}) = b_i^2 \sigma_v^2 + \tilde{\psi}_i$, where $\tilde{\psi}_i = \mathbb{E}(\psi_i \mid \mathcal{Z})$ is the smoothed variance (see the remark at the end of this section). The standardized error of the combined model is given by:

$$\varepsilon_i = \frac{\hat{\theta}_i - \boldsymbol{\beta}^\top \mathbf{z}_i}{\sqrt{b_i^2 \sigma_v^2 + \tilde{\psi}_i}}. \tag{2.4}$$

The direct estimate $\hat{\theta}_i$ provides information about θ_i. Rao and Molina (2015, Chapter 9, pages 271-272) give the conditional distribution of θ_i:

$$\theta_i \mid \mathcal{Z}, \hat{\theta}_i \sim \mathcal{N}\{\boldsymbol{\beta}^\top \mathbf{z}_i + \gamma_i(\hat{\theta}_i - \boldsymbol{\beta}^\top \mathbf{z}_i), (1-\gamma_i) b_i^2 \sigma_v^2\}, \tag{2.5}$$

where $\gamma_i = \frac{b_i^2 \sigma_v^2}{b_i^2 \sigma_v^2 + \tilde{\psi}_i}$.

The best predictor of θ_i, conditionally on $\hat{\theta}_i$ (Rao and Molina, 2015), is then given by:

$$\hat{\theta}_i^B = \mathbb{E}(\theta_i \mid \mathcal{Z}, \hat{\theta}_i) = \gamma_i \, \hat{\theta}_i + (1-\gamma_i) \, \boldsymbol{\beta}^\top \mathbf{z}_i. \tag{2.6}$$

In the remainder of this paper, the best predictor $\hat{\theta}_i^B$ will be called the B estimator.

In Sections 3 and 4, the theory is developed assuming that $\boldsymbol{\beta}$, σ_v^2 and $\tilde{\psi}_i$ are known. In Section 5, the estimation of these three quantities is discussed, which allows us to obtain an empirical version of the best predictor and our diagnostics.

Remark: In the literature on small area estimation, the theory is usually developed under the assumption that ψ_i is fixed. Therefore, it is implicitly assumed that $\tilde{\psi}_i = \psi_i$. When making inferences under the Fay-Herriot model, ψ_i cannot be expected to be fixed. For example, consider the case where θ_i is a proportion in the domain i and a stratified simple random sampling with replacement design is used with strata that coincide with domains. The direct estimator $\hat{\theta}_i$ is simply the sample proportion in the domain i and it is well known that its variance is given by $\psi_i = n_i^{-1} \theta_i (1-\theta_i)$. In this case, it is obvious that ψ_i is random since it depends on θ_i. It is also easy to show that $\tilde{\psi}_i = n_i^{-1}(\boldsymbol{\beta}^\top \mathbf{z}_i (1 - \boldsymbol{\beta}^\top \mathbf{z}_i) - b_i^2 \sigma_v^2) \neq \psi_i$ unless $v_i = \sigma_v = 0$. In the rest of this paper, the entire theory is developed under the usual assumption that $\tilde{\psi}_i = \psi_i$. In practice, these two variances are unknown and have to be estimated. Section 5 discusses the estimation of $\tilde{\psi}_i$ using a smoothing model. It can easily be shown that if a model-unbiased estimator, $\hat{\tilde{\psi}}_i$, is available, that is $\mathbb{E}(\hat{\tilde{\psi}}_i \mid \mathcal{Z}) = \tilde{\psi}_i$, then this estimator is also model-unbiased for ψ_i, that is $\mathbb{E}(\hat{\tilde{\psi}}_i - \psi_i \mid \mathcal{Z}) = 0$. The reverse is also true: a model-unbiased estimator for ψ_i will also be model-unbiased for $\tilde{\psi}_i$. Therefore, although $\tilde{\psi}_i \neq \psi_i$, both variances can be estimated by the same estimator. This suggests that the assumption $\tilde{\psi}_i = \psi_i$ may not be so critical in practice.

3. The mean square errors of the direct and B estimators

A mean square error criterion is often chosen to assess the efficiency of the B estimator given in equation (2.6). There are two natural possibilities: either consider the design MSE, or consider the model MSE (MSE with respect to the combined model 2.3).

The model MSE of the direct estimator $\hat{\theta}_i$ is:

$$\mathrm{MSE}_m(\hat{\theta}_i) = \mathbb{E}\{(\hat{\theta}_i - \theta_i)^2 \mid \mathcal{Z}\} = \tilde{\psi}_i$$

and the model MSE of the B estimator is:

$$\mathrm{MSE}_m(\hat{\theta}_i^B) = \mathbb{E}\{(\hat{\theta}_i^B - \theta_i)^2 \mid \mathcal{Z}\} = \gamma_i \tilde{\psi}_i.$$

The B estimator is thus always more efficient than the direct estimator with model-based inferences. This property is the result of the actual construction of the B estimator. On the other hand, and this is a legitimate question, is the B estimator always more efficient than the direct estimator under design-based inferences?

The design mean square errors of the direct and B estimators for the domain i are:

$$\begin{aligned}\operatorname{MSE}_p(\hat{\theta}_i) = \mathbb{E}\{(\hat{\theta}_i - \theta_i)^2 \mid \Omega\} &= \psi_i \\ &= \tilde{\psi}_i \end{aligned} \tag{3.1}$$

and

$$\begin{aligned}\operatorname{MSE}_p(\hat{\theta}_i^B) = \mathbb{E}\{(\hat{\theta}_i^B - \theta_i)^2 \mid \Omega\} &= \gamma_i^2 \psi_i + (1-\gamma_i)^2\, b_i^2 v_i^2 \\ &= \gamma_i \tilde{\psi}_i + (1-\gamma_i)^2 b_i^2 (v_i^2 - \sigma_v^2). \end{aligned} \tag{3.2}$$

Note that the second equality of (3.1) and (3.2) results from the assumption $\tilde{\psi}_i = \psi_i$. We observe that $\operatorname{MSE}_p(\hat{\theta}_i^B)$ can be very different from $\operatorname{MSE}_m(\hat{\theta}_i^B)$ when the unknown value v_i^2 is far from σ_v^2. Therefore, for a domain with a large value of v_i^2, $\operatorname{MSE}_m(\hat{\theta}_i^B)$ could be significantly smaller than $\operatorname{MSE}_p(\hat{\theta}_i^B)$ and lead to an inaccurate conclusion about the relative efficiency of the direct and B estimators.

By noticing that $\gamma_i \tilde{\psi}_i = (1-\gamma_i) b_i^2 \sigma_v^2$, we can show that $\operatorname{MSE}_p(\hat{\theta}_i^B) \le \operatorname{MSE}_p(\hat{\theta}_i)$ if and only if

$$v_i \in [-v_{L,i}; v_{L,i}],$$

where $v_{L,i} = \sigma_v \sqrt{\frac{1+\gamma_i}{\gamma_i}}$. Figure 3.1 shows the limit values $v_{L,i}/\sigma_v$ and $-v_{L,i}/\sigma_v$ as a function of γ_i. We note that when $|v_i| \le \sigma_v \sqrt{2}$, $\operatorname{MSE}_p(\hat{\theta}_i^B) \le \operatorname{MSE}_p(\hat{\theta}_i)$ for every value of γ_i. We also note that the direct estimator may become more efficient than the B estimator for domains where the local effect is large, especially when γ_i is not small. But how does one know if the local effect is large or not for a given domain i? This is the purpose of the following section where we present two diagnostics.

Figure 3.1 Limit values of the local effect standardized by σ_v versus γ_i.

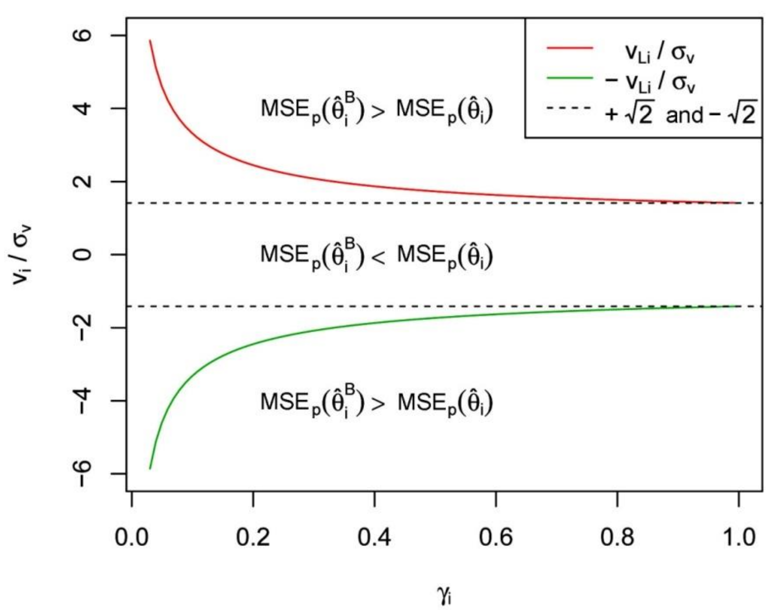

4. Two diagnostics to evaluate the local performance of the B estimator

4.1 An approach conditional on $\hat{\theta}_i$

From expression (2.5) in Section 2 and noting that $\gamma_i\left(\hat{\theta}_i - \boldsymbol{\beta}^\top \mathbf{z}_i\right) = b_i \sigma_v \sqrt{\gamma_i}\,\varepsilon_i,$ we obtain the conditional distribution of v_i:

$$v_i \,|\, \mathcal{Z}, \hat{\theta}_i \sim \mathcal{N}\left(\sigma_v \sqrt{\gamma_i}\,\varepsilon_i,\ (1-\gamma_i)\,\sigma_v^2\right).$$

Conditioning on $\hat{\theta}_i$ gives a better idea of the possible values v_i can take. In particular, when the value of γ_i is strictly greater than 0, the conditional distribution of v_i may deviate significantly from its unconditional distribution: $v_i \,|\, \mathcal{Z} \sim \mathcal{N}(0, \sigma_v^2).$

The first diagnostic is defined as the conditional probability:

$$\begin{aligned} D_{1i} &= \operatorname{Prob}\left(\operatorname{MSE}_p(\hat{\theta}_i^B) \le \operatorname{MSE}_p(\hat{\theta}_i) \,\middle|\, \mathcal{Z}, \hat{\theta}_i\right) \\ &= \operatorname{Prob}\left(-v_{L,i} \le v_i \le v_{L,i} \,\middle|\, \mathcal{Z}, \hat{\theta}_i\right). \end{aligned} \tag{4.1}$$

This diagnostic can be written as a function of γ_i and the standardized error (2.4):

$$\begin{aligned} D_{1i} = D_{1i}\left(\gamma_i, |\varepsilon_i|\right) &= \Phi\left\{\sqrt{\frac{\gamma_i}{1-\gamma_i}}\left(|\varepsilon_i| + \frac{\sqrt{1+\gamma_i}}{\gamma_i}\right)\right\} \\ &\quad - \Phi\left\{\sqrt{\frac{\gamma_i}{1-\gamma_i}}\left(|\varepsilon_i| - \frac{\sqrt{1+\gamma_i}}{\gamma_i}\right)\right\}, \end{aligned} \tag{4.2}$$

where $\Phi(\cdot)$ is the distribution function of the standard normal distribution. The proof of result (4.2) is given in Appendix A.

When this diagnostic takes values close to 0, we may conclude that $|v_i|$ is most likely larger than $v_{L,i}$ and that the direct estimator is preferable to the B estimator. To obtain a decision rule associated with this diagnostic, it is necessary to choose a threshold below which we decide to choose the direct estimator and above which the B estimator is chosen. A 50% threshold seems quite natural. Another idea is to apply an empirical approach and identify a break in the distribution of the values of diagnostic D_{1i} for the m domains.

This diagnostic is not entirely design-based because it involves the conditional distribution $v_i \,|\, \mathcal{Z}, \hat{\theta}_i$. It is therefore necessary to validate carefully the Fay-Herriot model before using it. Unfortunately, it is not possible to validate the assumptions on both v_i and e_i because the values of the parameters $\theta_i, i = 1, \ldots, m$ are not observed. However, the combined Fay-Herriot model (2.3) can be validated using model residuals (see, for example, Hidiroglou, Beaumont and Yung, 2019). These residuals are obtained by replacing the unknown quantities in the standardized error (2.4) with their estimates (see Section 5). A graph of residuals versus model predicted values is often suggested to validate the linearity assumption of the model. The normality assumption of the error $b_i v_i + e_i$ can be verified by a Q-Q plot of the residuals or normality tests such as the Shapiro-Wilk test. In case the model is not completely satisfactory, a conservative threshold of 75% may be appropriate.

The diagnostic in the following section is entirely design-based. It is therefore not dependent on the validity of the linking model. In this sense, it is considered more robust than the diagnostic (4.2). However, it relies on assumptions about the sampling errors e_i, discussed in Section 2, including the normality assumption of e_i.

4.2 Use of a design-based hypothesis test on the parameter v_i

In the design-based approach to inference, v_i is fixed and the standardized error (2.4) follows the distribution:

$$\varepsilon_i \mid \Omega \sim \mathcal{N}\left(v_i \frac{\sqrt{\gamma_i}}{\sigma_v},\ (1-\gamma_i)\right). \tag{4.3}$$

We have a unique observation of this random variable. We use it to test if $|v_i|$ is larger than $v_{L,i}$. We consider the test:

$$\mathbf{H_0}: |v_i| = v_{L,i} \quad \text{versus} \quad \mathbf{H_1}: |v_i| > v_{L,i}.$$

We use $|\varepsilon_i|$ as our test statistic. We expect that $|\varepsilon_i|$ will have smaller values under H_0 than under H_1. Let $\varepsilon_{\text{obs},i}$ be the observed value of the statistic ε_i and $P_i(v_i) = \text{Prob}\left(|\varepsilon_i| > |\varepsilon_{\text{obs},i}| \mid \Omega; v_i\right)$. The p-value of the test is defined as the probability that the statistic $|\varepsilon_i|$ is greater than the observed value $|\varepsilon_{\text{obs},i}|$ under the null hypothesis. Appendix B shows that the p-value is:

$$P_i(v_{L,i}) = P_i(-v_{L,i}) = \Phi(-\tau_i) + \Phi\left(-\tau_i - 2\frac{\sqrt{1+\gamma_i}}{\sqrt{1-\gamma_i}}\right),$$

where

$$\tau_i = \frac{|\varepsilon_{\text{obs},i}| - \sqrt{1+\gamma_i}}{\sqrt{1-\gamma_i}}.$$

Since the second term is often negligible compared to the first term, especially when $\tau_i > 0$ or γ_i is large, our second diagnostic is:

$$D_{2i} = D_{2i}\left(\gamma_i, |\varepsilon_{\text{obs},i}|\right) = \Phi\left(\frac{\sqrt{1+\gamma_i} - |\varepsilon_{\text{obs},i}|}{\sqrt{1-\gamma_i}}\right). \tag{4.4}$$

This second diagnostic can be interpreted as follows: When D_{2i} is small, we can assume that $|v_i|$ is likely to be larger than $v_{L,i}$ and the direct estimator is then preferred to the B estimator. For the choice of a decision threshold, values typically used as levels for hypothesis testing (e.g., 5% or 10%) can be used as a guide. With these small values, the B estimator is favoured. As with the previous diagnostic, the threshold can be determined by locating a break in the distribution of the values of diagnostic D_{2i} for the m domains.

4.3 Some properties of diagnostics 1 and 2

In this section, we study the behaviour of the functions $D_{1i}\left(\gamma_i, |\varepsilon_i|\right)$ and $D_{2i}\left(\gamma_i, |\varepsilon_i|\right)$ for limiting cases of γ_i and $|\varepsilon_i|$ and note their similarities and differences.

Case 1: $0 < \gamma_i < 1$ is fixed and $|\varepsilon_i| \to \infty$.

From equations (4.2) and (4.4) it can be shown that, for $|\varepsilon_i| > 0$, the two functions $D_{1i}\left(\gamma_i, |\varepsilon_i|\right)$ and $D_{2i}\left(\gamma_i, |\varepsilon_i|\right)$ decrease as $|\varepsilon_i|$ increases. In other words, the derivative of these functions with respect to $|\varepsilon_i|$ is negative. In addition, the limit when $|\varepsilon_i| \to \infty$ of these two functions tends toward 0. For a sufficiently large value of $|\varepsilon_i|$, the two diagnostics will therefore favour the direct estimator.

Case 2: $0 < \gamma_i < 1$ is fixed and $|\varepsilon_i| = 0$.

From equation (4.2), we observe that

$$D_{1i}\left(\gamma_i, 0\right) = \Phi\left(\sqrt{\frac{1+\gamma_i}{\gamma_i\left(1-\gamma_i\right)}}\right) - \Phi\left(-\sqrt{\frac{1+\gamma_i}{\gamma_i\left(1-\gamma_i\right)}}\right).$$

We can show that $D_{1i}\left(\gamma_i, 0\right)$ is minimized when $\gamma_i = -1+\sqrt{2}$. Therefore, $D_{1i}\left(\gamma_i, 0\right) \geq D_{1i}\left(-1+\sqrt{2}, 0\right) =$ 0.98. Since this value is close to 1, diagnostic 1 leads to choosing the B estimator in this case if a threshold of 0.50 or even 0.75 is chosen.

From equation (4.4) we obtain:

$$D_{2i}\left(\gamma_i, 0\right) = \Phi\left(\sqrt{\frac{1+\gamma_i}{1-\gamma_i}}\right).$$

We can show that, for $0 \leq \gamma_i < 1$, the function $D_{2i}\left(\gamma_i, 0\right)$ is minimized when $\gamma_i = 0$. Hence, $D_{2i}\left(\gamma_i, 0\right) \geq D_{2i}\left(0, 0\right) =$ 0.84. With a threshold smaller than 0.50, diagnostic 2 leads to the same decision as diagnostic 1 in this case, i.e. to choose the B estimator.

Case 3: $|\varepsilon_i| < \sqrt{2}$ is fixed and $\gamma_i \to 1$.

The two functions $D_{1i}\left(\gamma_i, |\varepsilon_i|\right)$ and $D_{2i}\left(\gamma_i, |\varepsilon_i|\right)$ tend toward 1 in this case. Therefore, diagnostics 1 and 2 lead to choosing the B estimator.

Case 4: $|\varepsilon_i| > \sqrt{2}$ is fixed and $\gamma_i \to 1$.

The two functions $D_{1i}\left(\gamma_i, |\varepsilon_i|\right)$ and $D_{2i}\left(\gamma_i, |\varepsilon_i|\right)$ tend toward 0 in this case. Diagnostics 1 and 2 lead here to choosing the direct estimator.

Case 5: $|\varepsilon_i|$ is fixed and $\gamma_i \to 0$.

The function $D_{1i}\left(\gamma_i, |\varepsilon_i|\right)$ tends toward 1 for any fixed value of $|\varepsilon_i|$. Therefore, diagnostic 1 favours the B estimator for small values of γ_i.

We note that $D_{2i}(0, |\varepsilon_i|) = \Phi(1 - |\varepsilon_i|)$. Therefore, contrary to Diagnostic 1, Diagnostic 2 will lead to choosing the direct estimator if $|\varepsilon_i|$ is sufficiently large even when γ_i is infinitely close to 0. For example, with a decision threshold at 0.05 and $\gamma_i = 0$, Diagnostic 2 favours the direct estimator when $|\varepsilon_i| > 1 - \Phi^{-1}(0.05) =$ 2.64.

In the first four cases above, both diagnostics lead to the same decision. There is a difference only in Case 5 where $\gamma_i \to 0$. We therefore expect that Diagnostic 2 will choose the direct estimator more often than Diagnostic 1 for small values of γ_i. Consider, for example, a threshold of 0.5 for Diagnostic 1 and of 0.05 for Diagnostic 2. For a threshold of 0.5, we can show that Diagnostic 1 leads to choosing the direct estimator as soon as $|\varepsilon_i|$ is larger than a value approximately equal to $\frac{\sqrt{1+\gamma_i}}{\gamma_i}$, i.e. as soon as $|\varepsilon_i| \gtrsim \frac{\sqrt{1+\gamma_i}}{\gamma_i}$. As for Diagnostic 2, for a threshold of 0.05, it leads to choosing the direct estimator as soon as $|\varepsilon_i| > \sqrt{1+\gamma_i} - \sqrt{1-\gamma_i}\, \Phi^{-1}(0.05)$. For $\gamma_i =$ 0.01, Diagnostic 1 thus leads to choosing the direct estimator when $|\varepsilon_i| \gtrsim$ 100.5, while Diagnostic 2 leads to choosing the direct estimator when $|\varepsilon_i| >$ 2.64. The gap narrows as γ_i increases. For example, for $\gamma_i =$ 0.2, Diagnostic 1 chooses the direct estimator when $|\varepsilon_i| \gtrsim$ 5.48 and Diagnostic 2 chooses the direct estimator when $|\varepsilon_i| >$ 2.57. The above discussion seems to suggest that Diagnostic 2 leads to choosing the direct estimator more often than Diagnostic 1. However, there are cases where Diagnostic 1 chooses the direct estimator contrary to Diagnostic 2. These cases generally occur for fairly large values of γ_i. For example, for $\gamma_i =$ 0.8, Diagnostic 1 chooses the direct estimator when $|\varepsilon_i| \gtrsim$ 1.68, while Diagnostic 2 chooses the direct estimator only when $|\varepsilon_i| >$ 2.08.

5. Empirical version of the B estimator and diagnostics

The theory has been developed assuming the parameters $\boldsymbol{\beta}$, σ_v^2 and $\tilde{\psi}_i$ are known. In practice, these quantities are unknown and the best predictor $\hat{\theta}_i^B$ cannot be used. They can be replaced by estimators $\hat{\boldsymbol{\beta}}$, $\hat{\sigma}_v^2$ and $\hat{\tilde{\psi}}_i$ to obtain the empirical best predictor (EB estimator):

$$\hat{\theta}_i^{\mathrm{EB}} = \hat{\gamma}_i\, \hat{\theta}_i + (1 - \hat{\gamma}_i)\, \hat{\boldsymbol{\beta}}^{\top} \mathbf{z}_i,$$

where $\hat{\gamma}_i = \frac{b_i^2 \hat{\sigma}_v^2}{b_i^2 \hat{\sigma}_v^2 + \hat{\tilde{\psi}}_i}$.

In what follows, we first discuss the estimation of $\boldsymbol{\beta}$ assuming σ_v^2 and $\tilde{\boldsymbol{\psi}} = (\tilde{\psi}_1, \ldots, \tilde{\psi}_m)^{\top}$ are known. This yields the estimator $\tilde{\boldsymbol{\beta}}(\sigma_v^2, \tilde{\boldsymbol{\psi}})$ of $\boldsymbol{\beta}$. Next, the estimation of σ_v^2 is discussed assuming that $\tilde{\boldsymbol{\psi}}$ is known and we obtain the estimator $\tilde{\sigma}_v^2(\tilde{\boldsymbol{\psi}})$ of σ_v^2. Finally, the estimation of the smoothed variances $\tilde{\psi}_i, i = 1, \ldots, m,$ is discussed. We denote the resulting estimators by $\hat{\tilde{\psi}}_i, i = 1, \ldots, m,$ and we let $\hat{\tilde{\boldsymbol{\psi}}} = (\hat{\tilde{\psi}}_1, \ldots, \hat{\tilde{\psi}}_m)^{\top}$. In practice, the smoothed variances must first be estimated and then successively we compute $\hat{\sigma}_v^2 = \tilde{\sigma}_v^2(\hat{\tilde{\boldsymbol{\psi}}})$ and $\hat{\boldsymbol{\beta}} = \tilde{\boldsymbol{\beta}}(\hat{\sigma}_v^2, \hat{\tilde{\boldsymbol{\psi}}})$, the estimates of σ_v^2 and $\boldsymbol{\beta}$.

Assuming σ_v^2 and $\tilde{\boldsymbol{\psi}} = (\tilde{\psi}_1, \ldots, \tilde{\psi}_m)^{\top}$ are known, the estimation of $\boldsymbol{\beta}$ can be done using the generalized least squares method, which is equivalent to the maximum likelihood estimation method under the assumption of independence and normality of the errors $b_i v_i + e_i$. We obtain:

$$\tilde{\boldsymbol{\beta}}\left(\sigma_v^2, \tilde{\boldsymbol{\psi}}\right) = \left(\sum_{i=1}^{m} \frac{\mathbf{z}_i \mathbf{z}_i^\top}{b_i^2 \sigma_v^2 + \tilde{\psi}_i}\right)^{-1} \sum_{i=1}^{m} \frac{\mathbf{z}_i \hat{\theta}_i}{b_i^2 \sigma_v^2 + \tilde{\psi}_i}.$$

Different methods exist for estimating σ_v^2. For example, the method of moments of Fay and Herriot (1979), the maximum likelihood or restricted maximum likelihood method can be used. The latter is more common in practice. All these methods consist of iteratively solving an estimation equation of the form $g\left(\sigma_v^2, \tilde{\boldsymbol{\psi}}\right) = 0$, where the function g depends on the method. The resulting estimator is denoted by $\tilde{\sigma}_v^2\left(\tilde{\boldsymbol{\psi}}\right)$. Rao and Molina (2015, Chapters 5 and 6) provide more details on the estimation of $\boldsymbol{\beta}$ and σ_v^2 and on the properties of estimators such as model consistency.

Before estimating σ_v^2 and $\boldsymbol{\beta}$ by $\hat{\sigma}_v^2 = \tilde{\sigma}_v^2\left(\hat{\tilde{\boldsymbol{\psi}}}\right)$ and $\hat{\boldsymbol{\beta}} = \tilde{\boldsymbol{\beta}}\left(\hat{\sigma}_v^2, \hat{\tilde{\boldsymbol{\psi}}}\right)$, it is first necessary to estimate the smoothed variance $\tilde{\psi}_i = \mathbb{E}(\psi_i \mid \mathcal{Z}), i = 1, \ldots, m$. We suppose that a design-unbiased estimator, $\hat{\psi}_i$, is available, i.e. $\mathbb{E}(\hat{\psi}_i \mid \Omega) = \psi_i$. Under this assumption, we observe that $\mathbb{E}(\hat{\psi}_i \mid \mathcal{Z}) = \tilde{\psi}_i$. The estimator $\hat{\psi}_i$ is therefore unbiased for the smoothed variance $\tilde{\psi}_i$ but can be very unstable when n_i is small. In general, it is preferable to model $\hat{\psi}_i$ given $\mathbf{z}_i$ to increase stability. The following smoothing model is frequently used in practice:

$$\log(\hat{\psi}_i) = \boldsymbol{\alpha}^\top \mathbf{x}_i + \eta_i,$$

where $\mathbf{x}_i$ is a function of $\mathbf{z}_i$, $\boldsymbol{\alpha}$ is a vector of model parameters and $\eta_i, i = 1, \ldots, m$, are independent and identically distributed errors with a mean equal to 0 and a variance equal to σ_η^2. It can easily be shown that

$$\tilde{\psi}_i = \mathbb{E}(\hat{\psi}_i \mid \mathcal{Z}) = \exp(\boldsymbol{\alpha}^\top \mathbf{x}_i)\, \Delta,$$

where $\Delta = \mathbb{E}\{\exp(\eta)\}$ and η is a random variable that follows the same distribution as the error term in the above smoothing model. A model-consistent estimator of $\boldsymbol{\alpha}$, denoted by $\hat{\boldsymbol{\alpha}}$, is obtained using the least squares method. Hidiroglou, Beaumont and Yung (2019) suggest estimating Δ by a model-consistent estimator, $\hat{\Delta}$, using a method of moments. The smoothed variance estimator is written as follows:

$$\hat{\tilde{\psi}}_i = \exp(\hat{\boldsymbol{\alpha}}^\top \mathbf{x}_i)\, \hat{\Delta},$$

where

$$\hat{\Delta} = \frac{\sum_{i=1}^{m} \hat{\psi}_i}{\sum_{i=1}^{m} \exp(\hat{\boldsymbol{\alpha}}^\top \mathbf{x}_i)}.$$

It can be expected that the design MSE of the EB estimator,

$$\mathrm{MSE}_p(\hat{\theta}_i^{\mathrm{EB}}) = \mathbb{E}\left\{(\hat{\theta}_i^{\mathrm{EB}} - \theta_i)^2 \mid \Omega\right\},$$

is greater than the design MSE of the B estimator given in equation (3.2). As mentioned above, the estimators of the parameters $\boldsymbol{\beta}$, σ_v^2, $\boldsymbol{\alpha}$ and Δ are model-consistent, as m increases, provided certain regularity conditions hold. Note also that that the design mean square error of the B estimator (see equation 3.2) does not depend on m. Therefore, the increase in the mean square error resulting from the estimation of these parameters can be expected to be modest when the number of domains is large. This suggests that, if m is large, the derivation of the bound $v_{L,i}$ will be little affected by the estimation of $\boldsymbol{\beta}$, σ_v^2, $\boldsymbol{\alpha}$ and Δ. Thus, our two diagnostics (4.2) and (4.4) should remain relevant even if the EB estimator is used instead of the B estimator. However, γ_i must be replaced by $\hat{\gamma}_i$ and ε_i by

$$\hat{\varepsilon}_i = \frac{\hat{\theta}_i - \hat{\boldsymbol{\beta}}^{\top}\mathbf{z}_i}{\sqrt{b_i^2\hat{\sigma}_v^2 + \hat{\tilde{\psi}}_i}}$$

in expressions (4.2) and (4.4) to be able to calculate these diagnostics with real data. As a result, we obtain $\hat{D}_{1i}$, the estimator of D_{1i}, and $\hat{D}_{2i}$, the estimator of D_{2i}.

6. Simulation study

A simulation study was conducted to evaluate the effectiveness of $\hat{D}_{1i}$ and $\hat{D}_{2i}$ in detecting which of the direct and EB estimators is preferable. We considered $m =$ 140 domains representing Canadian cities. In this simulation study, the vector of auxiliary variables is: $\mathbf{z}_i^{\top} = (1, z_{1i})$. The auxiliary variable z_{1i} is obtained from administrative files and is defined as the ratio of the number of employment insurance beneficiaries in city i to the number of people over 15 years of age in city i. The sample size in city i, n_i, was obtained from the Canadian Labour Force Survey (LFS). Of the 140 cities, 2 have a sample size smaller than 10, 10 have a sample size smaller than 30, 40 have a sample size smaller than 60, and 68 have a sample size smaller than 100, representing almost 50% of the cities. For these 68 cities, the estimated coefficients of variation of the LFS unemployment rates are in most cases too large to publish direct estimates of the unemployment rate; as a result, small area estimation techniques are required for these domains. In contrast, there are also 17 of the 140 cities with a sample size larger than 1,000 for which the direct estimate of the unemployment rate is reliable.

The population parameter θ_i was simulated for the m domains using the actual values of n_i and $\mathbf{z}_i$. It can be interpreted as the proportion of unemployed people in city i. The parameter θ_i was generated using the beta distribution with mean $\boldsymbol{\beta}^{\top}\mathbf{z}_i$ and variance σ_v^2, where $\sigma_v^2 = 7.58 \times 10^{-5}$ and $\boldsymbol{\beta}^{\top}$ = (0.0484, 0.95). These values of $\boldsymbol{\beta}$ and σ_v^2 were chosen from real data. We set $b_i = 1, i = 1, \ldots, m$. Then, we manually changed the values of θ_i for four domains (cities) in order to have a local effect v_i equal to $5\sigma_v$. Cities with different sample sizes were chosen: 10, 100, 501 and 3,773. In the rest of this section, the smallest of these four cities is identified by City 1 $(n_i = 10)$, the second smallest by City 2 $(n_i = 100)$, the second largest by City 3 $(n_i = 501)$ and the largest by City 4 ($n_i =$ 3,773).

A stratified simple random sampling with replacement design was considered where strata coincide with domains. The direct estimator $\hat{\theta}_i$ of θ_i is simply the proportion of sampled people in area i who have the characteristic of interest (e.g., being unemployed). Under such a simple design, it is easy to see that the direct estimator can be generated as follows: $\hat{\theta}_i = n_i^{-1}$ Binomial(n_i, θ_i). It is therefore not necessary to create the population of people in domain i to generate $\hat{\theta}_i$. We proceeded in this way in the simulation. The design variance of $\hat{\theta}_i$ is given by $\psi_i = n_i^{-1}\theta_i(1-\theta_i)$ and its estimator by $\hat{\psi}_i = (n_i - 1)^{-1}\hat{\theta}_i(1-\hat{\theta}_i)$. The smoothed variance $\tilde{\psi}_i$ is estimated using the smoothing model in Section 5 with $\mathbf{x}_i^\top = (1, \log(z_{1i}), \log(1-z_{1i}), \log(n_i))$.

In order to simulate a realistic scenario, the underlying assumptions of the Fay-Herriot model are not entirely satisfied in our simulation. For example, the errors v_i and e_i do not exactly follow normal distributions. We used a beta distribution to generate θ_i; the normality assumption of v_i is therefore not satisfied although the deviation from the normal distribution is not severe in our simulation. The estimates $\hat{\theta}_i$ were generated from a binomial distribution,which can be approximated by a normal distribution for domains with a large value of n_i. The relationship between the simulated estimates $\hat{\theta}_i$ and the auxiliary vectors $\mathbf{z}_i$ is similar to the one observed with the real LFS estimates. Moreover, our simulation scenario is such that the assumption $\tilde{\psi}_i = \psi_i$ is not satisfied since, for this simple design,

$$\tilde{\psi}_i = n_i^{-1}\left(\boldsymbol{\beta}^\top\mathbf{z}_i(1-\boldsymbol{\beta}^\top\mathbf{z}_i) - b_i^2\sigma_v^2\right)$$

(see remark in Section 2). However, we note that the correlation coefficient between $\tilde{\psi}_i$ and ψ_i is 0.98, which indicates that the deviation from the assumption $\tilde{\psi}_i = \psi_i$ is modest. As mentioned in the previous paragraph, the smoothing model in Section 5 is used to estimate $\tilde{\psi}_i$. This allows us to remain in a realistic framework where the postulated smoothing model is different from the true model used to generate the estimates $\hat{\psi}_i$.

We conducted a design-based simulation study, i.e., the population parameters $\theta_i, i = 1, \ldots, m,$ were generated only once. We repeated sample selection K = 10,000 times. For each replicate k, $k = 1, \ldots, K$, a direct estimate $\hat{\theta}_i(k)$ was generated and a smoothed variance estimate $\hat{\tilde{\psi}}_i(k)$ was calculated as described above. The EB estimate was then calculated as:

$$\hat{\theta}_i^{\mathrm{EB}}(k) = \hat{\gamma}_i(k)\,\hat{\theta}_i(k) + (1-\hat{\gamma}_i(k))\,\hat{\boldsymbol{\beta}}(k)^\top\mathbf{z}_i,$$

where $\hat{\gamma}_i(k) = \frac{\hat{\sigma}_v^2(k)}{\hat{\sigma}_v^2(k) + \hat{\tilde{\psi}}_i(k)}$ and $\hat{\boldsymbol{\beta}}(k)$ and $\hat{\sigma}_v^2(k)$ are calculated as described in Section 5. The generalized least squares method was used to obtain $\hat{\boldsymbol{\beta}}(k)$ and the restricted maximum likelihood method was used to obtain $\hat{\sigma}_v^2(k)$. Calculations were performed using Statistics Canada's small area estimation system (Hidiroglou, Beaumont and Yung, 2019).

For each replicate, standardized residuals $\hat{\varepsilon}_i(k)$ and diagnostics $\hat{D}_{1i}(k)$ and $\hat{D}_{2i}(k)$ were also calculated for the m domains. We recorded whether the direct estimator was preferred over the EB estimator for each of the two diagnostics. Decision thresholds were used for this purpose. Below the thresholds, the direct estimator is used. For Diagnostic 1, thresholds of 50% and 75% were used and for Diagnostic 2, thresholds of 5% and 25% were used.

From the previous quantities, calculated for each of the 10,000 replicates, the Monte Carlo averages of Diagnostics 1 and 2 were calculated for the m domains: $\bar{\hat{D}}_{1i}$ and $\bar{\hat{D}}_{2i}$. The selection rate of the direct estimator was also calculated for each of the two diagnostics, i.e., the percentage of times a given diagnostic led to the selection of the direct estimator.

The Monte Carlo approximation of $\text{MSE}_p(\hat{\theta}_i^{\text{EB}})$ was calculated as:

$$\text{MSE}_{\text{MC}}(\hat{\theta}_i^{\text{EB}}) = \frac{1}{K}\sum_{k=1}^{K}\left(\hat{\theta}_i^{\text{EB}}(k) - \theta_i\right)^2.$$

From this Monte Carlo MSE, the relative efficiency of the EB estimator was calculated as:

$$\frac{\text{MSE}_{\text{MC}}(\hat{\theta}_i^{\text{EB}}) - \psi_i}{\psi_i}. \tag{6.1}$$

This ratio is positive when the EB estimator is less efficient than the direct estimator under the design. A diagnostic is potentially useful if it is negatively correlated with this ratio.

Figures 6.1 and 6.2 present the Monte Carlo averages of Diagnostic 1 and 2 respectively as a function of the relative efficiency of the EB estimator defined in equation (6.1). The four cities whose values of v_i have been changed, Cities 1 to 4, are shown in *purple*, *orange*, *green* and *red*. In the legend, the sample size of these cities has been indicated. The values of the parameter γ_i for Cities 1 to 4 are 0.01, 0.08, 0.35 and 0.81 respectively. All other cities are shown in blue.

First, we can see in Figures 6.1 and 6.2 that the EB estimator is more efficient than the direct estimator for City 1 (in purple) since this city is to the left of the vertical line (negative relative efficiency) despite the strong local effect. The explanation of this phenomenon is shown in Figure 3.1. It shows that the range of values of v_i for which the B estimator is more efficient than the direct estimator increases as γ_i decreases. Since γ_i is small for City 1 $(\gamma_i = 0.01)$, it is not surprising to observe a negative relative efficiency despite a pronounced local effect. For City 2 (in orange), the direct estimator is slightly more efficient than the EB estimator. On the other hand, for Cities 3 (in green) and 4 (in red), the direct estimator is much more efficient than the EB estimator. Note also that there are five cities for which the direct estimator is more efficient than the EB estimator: Cities 2 to 4 as well as two other cities whose values of v_i were randomly generated and not manually modified. One of these cities has the smallest value of v_i and the other has the largest value of v_i after excluding the four cities that had their value manually modified. These two cities have large values of γ_i (0.62 and 0.49).

Figures 6.1 and 6.2 indicate that our two diagnostics seem to be quite effective in detecting cases where the direct estimator is more efficient than the EB estimator except for City 2 $(n_i = 100)$ where the Monte Carlo average of Diagnostic 1 is very high at 0.97. However, this is a domain where choosing the least efficient estimator is not really problematic since there is very little difference between the efficiencies of the two estimators. Apart from this specific case, Diagnostic 1 seems to have better properties than Diagnostic 2. The Monte Carlo average of Diagnostic 1 is very close to 1 when the EB estimator is significantly more efficient than the direct estimator, decreases slowly when the efficiencies of the two estimators approach each other and becomes small when the direct estimator is significantly more efficient than the EB estimator. Not exactly the same behaviour is observed for Diagnostic 2. The Monte Carlo average of Diagnostic 2 is small when the direct estimator is significantly more efficient than the EB estimator but it is not close to 1 when the EB estimator is significantly more efficient than the direct estimator. Furthermore, it seems to increase when the efficiencies of the two estimators come closer, which is counterintuitive.

Figure 6.1 Monte Carlo average of Diagnostic 1 estimates for the 140 cities versus the relative efficiency of the EB estimator.

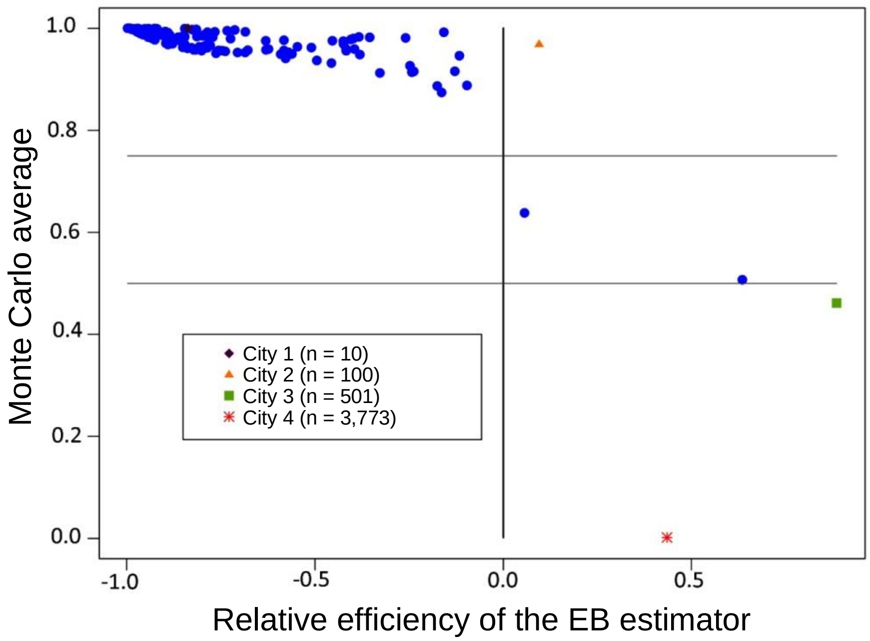

Figure 6.2 Monte Carlo average of Diagnostic 2 estimates for the 140 cities versus the relative efficiency of the EB estimator.

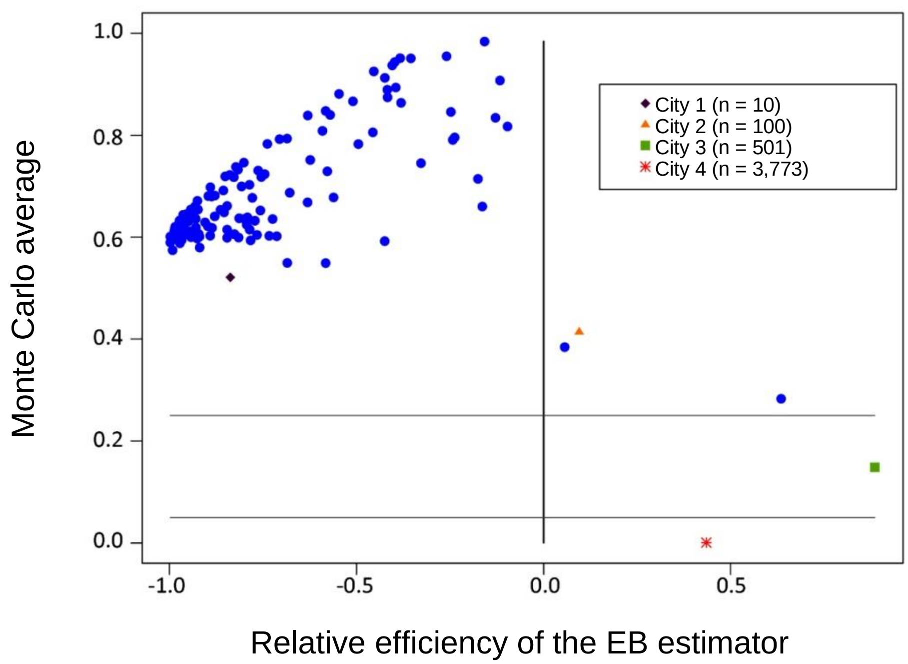

Figures 6.3 and 6.4 show the selection rate of the direct estimator over the 10,000 replicates for Diagnostics 1 and 2. Similar conclusions can be drawn as those obtained by analyzing Figures 6.1 and 6.2.

As expected, the thresholds of 75% for Diagnostic 1 and 25% for Diagnostic 2 allow better detection of cases where the direct estimator is more efficient than the EB estimator, but these thresholds also lead to the direct estimator being chosen a little too often when it was less efficient than the EB estimator. This is particularly notable for Diagnostic 2. This error can be dampened by decreasing the thresholds, but this also reduces the selection rate of the direct estimator when it is more efficient than the EB estimator. As noted earlier, Diagnostic 1 appears to have better properties than Diagnostic 2, regardless of the thresholds chosen, with very small selection rates of the direct estimator when it is significantly less efficient than the EB estimator. This seems to show the limitations of a fully design-based approach, such as the one presented in Section 4.2, to address the challenge of small domain sample sizes.

Figure 6.3 Direct estimator selection rate for Diagnostic 1.

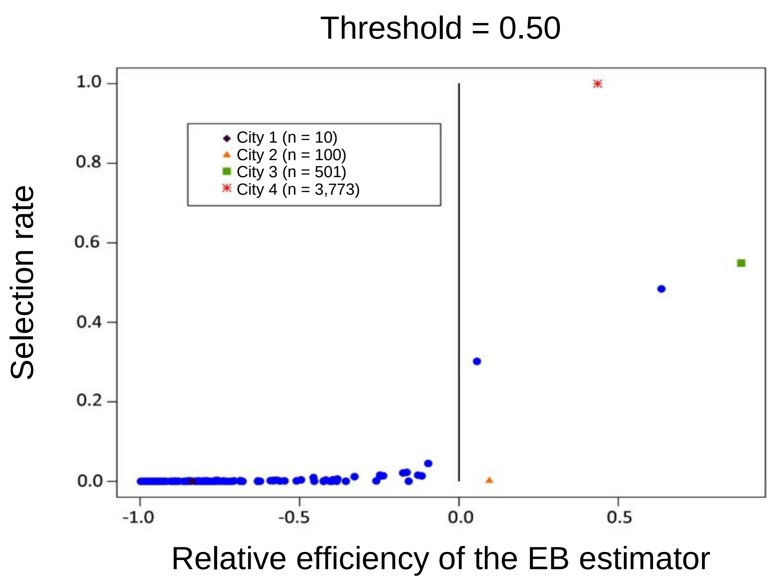

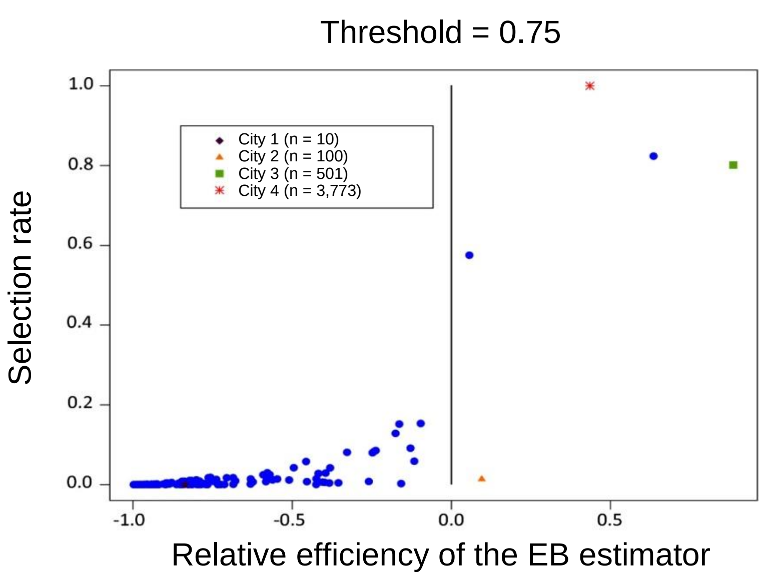

Figure 6.4 Direct estimator selection rate for Diagnostic 2.

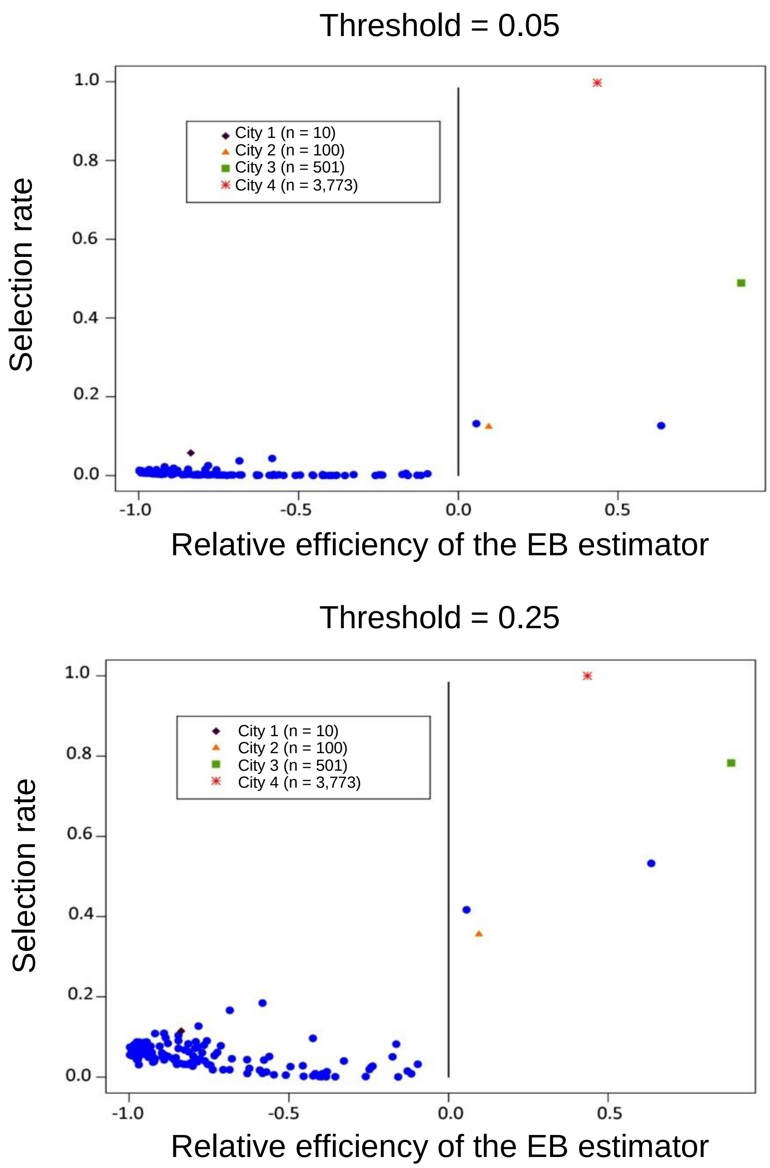

7. Conclusion

Users of small area estimates are usually interested in only one domain. Therefore, they seek a quality indicator that applies to their domain and not an overall indicator. The design MSE of small area estimators is a conceptually attractive quality indicator since it conditions on the unexplained local effect. However, it is known that design-unbiased estimators of the design MSE are generally unstable when the domain sample size is small. To circumvent this problem, we proposed two diagnostics that are intended to identify domains where the design MSE of the direct estimator is smaller than that of the EB estimator. Our simulation results seem promising and allow us to envision the implementation of a useful indicator for choosing between the direct and EB estimators for a particular domain. In future research, it would be interesting to evaluate the efficiency of a hybrid estimator that would leverage these diagnostics.

Appendix

A. Proof of equivalence between equations (4.1) and (4.2)

Using equation (4.1) and the conditional distribution of v_i given in Section 4.1, we have:

$$\begin{aligned} D_{1i} &= \text{Prob}\left(-v_{L,i} \le v_i \le v_{L,i} \middle| Z, \hat{\theta}_i\right) \\ &= \text{Prob}\left(\frac{-v_{L,i} - \sigma_v\sqrt{\gamma_i}\varepsilon_i}{\sigma_v\sqrt{1-\gamma_i}} \le \frac{v_i - \sigma_v\sqrt{\gamma_i}\varepsilon_i}{\sigma_v\sqrt{1-\gamma_i}} \le \frac{v_{L,i} - \sigma_v\sqrt{\gamma_i}\varepsilon_i}{\sigma_v\sqrt{1-\gamma_i}} \middle| Z, \hat{\theta}_i\right). \end{aligned}$$

Replacing $v_{L,i}$ with $\sigma_v \sqrt{\frac{1+\gamma_i}{\gamma_i}}$ results in:

$$\begin{aligned} D_{1i} &= \text{Prob}\left(\frac{-\sqrt{(1+\gamma_i)/\gamma_i} - \sqrt{\gamma_i}\varepsilon_i}{\sqrt{1-\gamma_i}} \le \frac{v_i - \sigma_v\sqrt{\gamma_i}\varepsilon_i}{\sigma_v\sqrt{1-\gamma_i}} \le \frac{\sqrt{(1+\gamma_i)/\gamma_i} - \sqrt{\gamma_i}\varepsilon_i}{\sqrt{1-\gamma_i}} \middle| Z, \hat{\theta}_i\right) \\ &= \Phi\left\{\sqrt{\frac{\gamma_i}{1-\gamma_i}}\left(-\varepsilon_i + \frac{\sqrt{1+\gamma_i}}{\gamma_i}\right)\right\} - \Phi\left\{\sqrt{\frac{\gamma_i}{1-\gamma_i}}\left(-\varepsilon_i - \sqrt{\frac{1+\gamma_i}{\gamma_i}}\right)\right\}. \end{aligned}$$

Since for any value t, we have $\Phi(t) = 1 - \Phi(-t)$ then

$$D_{1i} = \Phi\left\{\sqrt{\frac{\gamma_i}{1-\gamma_i}}\left(\varepsilon_i + \frac{\sqrt{1+\gamma_i}}{\gamma_i}\right)\right\} - \Phi\left\{\sqrt{\frac{\gamma_i}{1-\gamma_i}}\left(\varepsilon_i - \frac{\sqrt{1+\gamma_i}}{\gamma_i}\right)\right\}.$$

We notice that D_{1i} is a symmetric function of ε_i around 0, i.e. $D_{1i}(\varepsilon_i) = D_{1i}(-\varepsilon_i)$. Therefore, we can rewrite D_{1i} as in equation (4.2):

$$D_{1i} = \Phi\left\{\sqrt{\frac{\gamma_i}{1-\gamma_i}}\left(|\varepsilon_i| + \frac{\sqrt{1+\gamma_i}}{\gamma_i}\right)\right\} - \Phi\left\{\sqrt{\frac{\gamma_i}{1-\gamma_i}}\left(|\varepsilon_i| - \frac{\sqrt{1+\gamma_i}}{\gamma_i}\right)\right\}.$$

B. p-value associated with the test statistic $|\varepsilon_i|$

First, recall that $P_i(v_i) = \text{Prob}\left(|\varepsilon_i| > |\varepsilon_{\text{obs},i}| \mid \Omega; v_i\right)$. We define the p-value as the maximum of $P_i(v_{L,i})$ and $P_i(-v_{L,i})$. Since $\tau_i = \frac{|\varepsilon_{\text{obs},i}| - \sqrt{1+\gamma_i}}{\sqrt{1-\gamma_i}}$, we can then write:

$$\begin{aligned} P_i(v_i) &= \text{Prob}\left(|\varepsilon_i| > \sqrt{1+\gamma_i} + \sqrt{1-\gamma_i}\,\tau_i \middle| \Omega; v_i\right) \\ &= \text{Prob}\left(\varepsilon_i > \sqrt{1+\gamma_i} + \sqrt{1-\gamma_i}\,\tau_i \middle| \Omega; v_i\right) \\ &\quad + \text{Prob}\left(\varepsilon_i < -\sqrt{1+\gamma_i} - \sqrt{1-\gamma_i}\,\tau_i \middle| \Omega; v_i\right). \end{aligned}$$

Using the standardized error distribution (4.3), we obtain:

$$\begin{aligned} P_i(v_i) &= \text{Prob}\left(\frac{\varepsilon_i - v_i \sqrt{\gamma_i}/\sigma_v}{\sqrt{1-\gamma_i}} > \frac{\sqrt{1+\gamma_i} - v_i \sqrt{\gamma_i}/\sigma_v}{\sqrt{1-\gamma_i}} + \tau_i \,\middle|\, \Omega; v_i \right) \\ &\quad + \text{Prob}\left(\frac{\varepsilon_i - v_i \sqrt{\gamma_i}/\sigma_v}{\sqrt{1-\gamma_i}} < \frac{-\sqrt{1+\gamma_i} - v_i \sqrt{\gamma_i}/\sigma_v}{\sqrt{1-\gamma_i}} - \tau_i \,\middle|\, \Omega; v_i \right) \\ &= \Phi\left(\frac{-\sqrt{1+\gamma_i} + v_i \sqrt{\gamma_i}/\sigma_v}{\sqrt{1-\gamma_i}} - \tau_i \right) \\ &\quad + \Phi\left(\frac{-\sqrt{1+\gamma_i} - v_i \sqrt{\gamma_i}/\sigma_v}{\sqrt{1-\gamma_i}} - \tau_i \right). \end{aligned}$$

Using the expression $v_{L,i} = \sigma_v \sqrt{\frac{1+\gamma_i}{\gamma_i}}$, we have:

$$\begin{aligned} P_i(v_i) &= \Phi\left(-\tau_i + \frac{\sqrt{1+\gamma_i}}{\sqrt{1-\gamma_i}} \left[\frac{v_i}{v_{L,i}} - 1 \right] \right) \\ &\quad + \Phi\left(-\tau_i + \frac{\sqrt{1+\gamma_i}}{\sqrt{1-\gamma_i}} \left[-\frac{v_i}{v_{L,i}} - 1 \right] \right). \end{aligned}$$

Under the null hypothesis H_0, $v_i = v_{L,i}$ or $v_i = -v_{L,i}$ and in both cases the above equation reduces to:

$$P_i(v_{L,i}) = P_i(-v_{L,i}) = \Phi(-\tau_i) + \Phi\left(-\tau_i - 2\frac{\sqrt{1+\gamma_i}}{\sqrt{1-\gamma_i}} \right).$$

We will now show that if we reject H_0 (with a threshold smaller than 0.5 such as 0.1) then we would reject even more strongly the null hypothesis $H_0^*: |v_i| = v_i^*$ for any value $0 \le v_i^* < v_{L,i}$. First, if $\tau_i \le 0$, i.e., $|\varepsilon_{\text{obs},i}| \le \sqrt{(1+\gamma_i)}$, we observe that $P_i(v_{L,i}) = P_i(-v_{L,i}) \ge 0.5$ and we never reject the null hypothesis H_0. Second, if $\tau_i > 0$, we can easily show that the function $P_i(v_i)$ is increasing in v_i over the interval $[0, v_{L,i}]$. We also note that it is a function of v_i that is symmetrical around $v_i = 0$ since $P_i(v_i) = P_i(-v_i)$. Consequently, $P_i(v_i)$ is decreasingon the interval $[-v_{L,i}, 0]$, is minimum when $v_i = 0$ and maximum when $v_i = v_{L,i}$ and $v_i = -v_{L,i}$. Therefore, when $|v_i| < v_{L,i}$, we have:

$$P_i(v_i) < P_i(v_{L,i}) = \Phi(-\tau_i) + \Phi\left(-\tau_i - 2\frac{\sqrt{1+\gamma_i}}{\sqrt{1-\gamma_i}} \right).$$

References

Fay, R.E., and Herriot, R.A. (1979). Estimation of income from small places: An application of James-Stein procedures to census data. *Journal of the American Statistical Association*, 74, 269-277.

Hidiroglou, M.A., Beaumont, J.-F. and Yung, W. (2019). Development of a small area estimation system at Statistics Canada. *Survey Methodology*, 45, 1, 101-126. Paper available at https://www150.statcan.gc.ca/n1/en/pub/12-001-x/2019001/article/00009-eng.pdf.

Pfeffermann, D., and Ben-Hur, D. (2019). Estimation of randomisation mean square error in small area estimation. *International Statistical Review*, 87, S1, S31-S49.

Rao, J.N.K., and Molina, I. (2015). *Small Area Estimation*. John Wiley & Sons, Inc., Hoboken, New Jersey.

Rao, J.N.K., Rubin-Bleuer, S. and Estevao, V.M. (2018). Measuring uncertainty associated with model-based small area estimators. *Survey Methodology*, 44, 2, 151-166. Paper available at https://www150.statcan.gc.ca/n1/en/pub/12-001-x/2018002/article/54958-eng.pdf.

Rivest, L.-P., and Belmonte, E. (2000). A conditional mean squared error of small area estimators. *Survey Methodology*, 26, 1, 67-78. Paper available at https://www150.statcan.gc.ca/n1/en/pub/12-001-x/2000001/article/5179-eng.pdf.